

APPLICATION UNDER UNITED STATES PATENT LAWS

Atty. Dkt. No. PM 270174
(M#)

Invention: METHOD AND SYSTEM FOR A ROUTING MECHANISM TO SUPPORT TWO-WAY RSVP RESERVATIONS

Inventor (s): SEDDIGH, Nabil N.
NANDY, Biswajit B.

Pillsbury Madison & Sutro LLP
Intellectual Property Group
1100 New York Avenue, NW
Ninth Floor
Washington, DC 20005-3918
Attorneys
Telephone: (202) 861-3000

This is a:

- Provisional Application
- Regular Utility Application
- Continuing Application
 - The contents of the parent are incorporated by reference
- PCT National Phase Application
- Design Application
- Reissue Application
- Plant Application
- Substitute Specification
Sub. Spec Filed
in App. No. _____ / _____
- Marked up Specification re
Sub. Spec. filed
In App. No. _____ / _____

SPECIFICATION

METHOD AND SYSTEM FOR A ROUTING MECHANISM TO SUPPORT TWO-WAY RSVP RESERVATIONS

5

BACKGROUND OF THE INVENTION

1. Field of the Invention

10 Aspects of the invention generally relate to the field of network communications. Specifically, aspects of the present invention relate to a method and system for a quality of service mechanism that supports two-way RSVP reservations.

2. Description of Background Information

15 In an information age, achieving the highest network service quality is as important as developing best class of networking products. This is particularly so when new applications, such as Voice over IP (VoIP) and video conferencing, place new demands on the network. Various service models, network protocols, and standards have been proposed, aiming at improving the network management efficiency and maximizing the utilization of the network.

20 Quality of Service (QoS) mechanisms are proposed to provide the necessary level of service to applications and to maintain an expected quality level. The first concerted effort at providing QoS for IP networks focused on the Integrated Services (IntServ) architecture, which provides per-flow end-to-end QoS guarantees based on signaled requests from end-host applications. RSVP emerged as the signaling protocol of choice for IntServ.

25 The Differentiated Services (DiffServ or DS) architecture has recently become the preferred method that addresses QoS issues in IP networks. In DiffServ, individual flows are

grouped into aggregated traffic classes by edge devices such as edge routers. Packets are marked to reflect the treatment required by the traffic class. Core routers differentially treat packets according to the traffic class marking. Recently, RSVP has been used as the protocol of choice to enable dynamic signaling and admission control in DiffServ IP networks.

5 FIG. 1 shows a DS framework. A DS framework comprises a plurality of DS domains, each of which is a set of contiguous DS-compliant networks containing DS-compliant nodes. An end-to-end differentiated service is obtained by the concatenation of per-domain services and service level agreements between adjoining domains along a source-to-destination traffic path. The exemplary DS shown in FIG. 1 is a concatenation of 4 DS
10 domains, each of which has ingress devices (E₁, E₃, E₅, E₇), egress devices (E₂, E₄, E₆, E₈), or collectively referred to herein as edge devices, and core devices (C₁, C₂, C₃, C₄).

Per-domain services are realized by traffic conditioning at the edge and simple differentiated forwarding mechanisms at the core of the network. To build an end-to-end service, subscribed traffic profiles for customers are maintained by using traffic filters. The
15 traffic is metered and measured against the associated traffic profiles. Packets are grouped into a set of coarse aggregate flows that receive differentiated treatment at the network core.

In both IntServ and DiffServ architectures, RSVP is used for one-directional resource reservation. For an application that requires 2-way traffic on the network such as a Voice-over-IP application, two separate reservations need to be made. FIG. 2 depicts a typical receiver-driven RSVP signaling scheme in a DS framework. In FIG. 2, a 3-way handshake
20 RSVP is used to complete a one-directional reservation. The first handshake is from a sender 310 to a receiver 320 with a PATH message to probe a path. The second is from the receiver 320 to the sender 310 with a RESV message that initiates the reservation at edge devices
310.

along the path. The third is from the sender 310 back to the receiver 320 with an acknowledgement message. An RSVP request is processed only at edge devices and the reservation is made via the communication with the Policy and Decision Point (PDP) at each domain. With this scheme, a two-directional reservation requires two rounds of 3-way handshake, resulting in a total of six-way handshake.

SUMMARY OF THE INVENTION

A method and system, consistent with the principles of the present invention, provides support for two-way RSVP reservations. A first embodiment of the present invention is a 10 sender driven protocol that completes a bi-directional RSVP reservation for a two-way communication application in a 3-way handshake. A second embodiment of the present invention is a receiver driven protocol that completes a bi-directional RSVP reservation in a 4-way handshake.

15

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is further described in the detailed description which follows, with reference to the drawings by way of non-limiting embodiments of the present invention. It is noted that, throughout the description, like reference numerals represent similar parts of the present invention throughout the several views and wherein:

20 FIG. 1 illustrates an exemplary differentiated services network, connecting a sender and a receiver via four Differentiated Service domains;

FIG. 2 shows a known example of a receiver-driven RSVP reservation scheme in a DS network;

FIG. 3 illustrates an embodiment of the invention, in which a sender-driven RSVP scheme completes a two-directional reservation in a three-way handshake protocol;

FIG. 4 is a flowchart for the first party, in a two-way communication application, who initiates the two-way communication;

5 FIG. 5 is a flowchart for an ingress policy enforcement device of a DS domain;

FIG. 6 is a flowchart for an egress policy enforcement device of a DS domain;

FIG. 7 is flowchart for a PDP of a DS domain;

FIGs. 8 and 9 are a flowchart for the second party in a two-way communication application;

10 FIG. 10 illustrates a second embodiment of the invention, in which a receiver-driven RSVP scheme completes a two-directional reservation in a four-way handshake protocol;

FIGs. 11 through 13 present a flowchart for the first party in the second embodiment, who initiates the two-way communication;

FIG. 14 is a flowchart for an ingress policy enforcement device of a DS domain in the

15 second embodiment;

FIG. 15 is a flowchart for an egress policy enforcement device of a DS domain in the second embodiment; and

FIGs. 16 and 17 are a flowchart for the second party in the second embodiment.

20

DETAILED DESCRIPTION OF SEVERAL EMBODIMENTS

FIG. 3 illustrates a first embodiment of the invention, in which an RSVP scheme for reserving network resource for a two-way communication application is shown. The scheme

illustrated in FIG. 3 is sender-driven and it completes a two directional resource reservation in a 3-way handshake. In FIG. 3, there are two participants in a two-way communication. The party that initiates the handshake is identified herein as the sender (310). The other party is identified herein as the receiver (320). Even though either party may send information to the
5 other once a two-way communication is established, the terms sender and the receiver will be used throughout this description according to the definition above. Forward direction is herein defined as from the sender 310 to the receiver 320 and reverse direction is herein defined as from the receiver 320 to the sender 310.

There are four concatenated network domains in the exemplary network in FIG. 3.

- 10 Looking in the forward direction, E₁ and E₂ (330a and 340a) are the ingress and egress policy enforcement devices of the first domain, or collectively referred to herein as edge policy enforcement device. A policy enforcement device is used herein as a general term. A policy enforcement device may be implemented as a router or as a multiplexer. Similarly, E₃ and E₄ (330b, 340b) are the ingress and egress policy enforcement devices of the second domain, E₅
15 and E₆ (330c and 340c) are the ingress and egress policy enforcement devices of the third domain, and E₇ and E₈ (330d and 340d) are the ingress and egress policy enforcement devices of the fourth domain, respectively.

Even though ingress and egress policy enforcement devices are reversed when the traffic flow is in the reverse direction, for the clarity of this presentation, ingress and egress
20 policy enforcement devices are herein termed with respect to the forward direction. Each domain in FIG. 3 has its own Policy Decision Point (PDP) (350a, 350b, 350c, or 350d). To perform resource reservation, an edge policy enforcement device, either ingress or egress,

may communicate with its domain PDP using Common Open Policy Services – RSVP (COPS-RSVP) protocol.

In the 3-way handshake RSVP scheme described in FIG. 3, the first pass reserves the network resources in forward direction. In the second pass, network resources in reverse direction are reserved. The third pass simply delivers an acknowledgement message that completes the 3-way handshake.

In FIG. 3, sender 310 initiates a 3-way handshake by generating a PATH message. The PATH message contains information identifying the sending session as well as the traffic profile for the sender 310. At the ingress of each domain, the edge policy enforcement device

- 10 (330a, 30b, 330c, or 330d) intercepts the PATH message and performs policy and bandwidth admission control by communicating with the corresponding PDP (350a, 350b, 350c, or 350d) using COPS-RSVP. If the admission control decision is granted, the PATH message is passed on to the egress policy enforcement device (340a, 340b, 340c, or 340d) of the same domain. The ingress policy enforcement devices do not add their address to the NEXT_HOP
15 object in the PATH message because the RESV message in the second pass will not go through ingress policy enforcement devices. At the egress of each domain, the edge policy enforcement device (340a, 340b, 340c, or 340d) examines the PATH message and adds its own address to the NEXT_HOP object to ensure that the RESV message in the second pass will be sent to it.

- 20 When the PATH message reaches the receiver 320, the reservation in one direction (forward direction) is successful. In this case, receiver 320 generates an RESV message. Such generated RESV message does not simply echo the PATH message, as would be the case in a conventional method. Instead, it carries the network reservation information from the receiver

320 to the sender 310 (in the reverse direction). This RESV message is sent hop-by-hop from the receiver 320 to the sender 310 via the egress policy enforcement device (340a, 340b, 340c, 340d) of each domain, following the addresses specified in the NEXT_HOP object. At each egress policy enforcement device, communication with the PDP of the same domain decides whether the resource reservation request in the reverse direction will be admitted or not. If the request is admitted, the PDP will install necessary filters and traffic profiles via COPS-PROvisioning (COPS-PR). When RESV message reaches the sender 310, the reservation in the second direction (reverse direction) is successful. The sender 310 then sends a RESV_Confirm message directly to the receiver 320 to complete the 3-way handshake.

In the embodiment illustrated in FIG. 3, a two-way reservation is considered successful when the admission control decisions are granted in both directions. A failure in reserving the resource in the forward direction at any ingress policy enforcement device (330a, 330b, 330c...) may be signaled by sending a PATH_ERR message from that policy enforcement device to the sender 310. A failure in reserving the resource in the reverse direction at any egress policy enforcement device (340d, 340c, 340b...) may be signaled by sending an RESV_ERR message back to the receiver 320. To inform the sender 310 of a failure in reserving the resources in reverse direction, either a PATH_ERR message may be sent to the sender 310 or a time out mechanism may be applied at the sender 310.

FIG. 4 shows a flowchart for the sender 310. To initiate a two-directional RSVP reservation, sender 310 first constructs a PATH message at 410. This PATH message carries information identifying the sending session as well as the traffic profile for the sender 310, and initiates the resource reservation for forward direction. The PATH message is sent out at

420. Time may be marked at 430 so that a time reference for a time-out mechanism may be established. The sender 310 then enters a waiting mode for a return message.

If a message is not received before a time-out at 440, the handshake is aborted. If a message is received within the time-out at 440, the type of the message is determined. It is 5 first examined at 445 to see whether it is a PATH_ERR message. A PATH_ERR message indicates that the reservation for the forward direction has failed. In this case, the sender 310 aborts the 3-way handshake. If the return message is not a PATH_ERR message, it may be further examined at 450 to see whether it is an RESV message. If it is not an RESV message, the sender 310 goes back to 440 to wait for a return message. When the received message is 10 an RESV, it means that the reservation in both forward and reverse directions have been successful. In this case, the sender 310 constructs a third message, an RESV_Confirm message, at 470 and sends it out at 480 directly to the receiver 320 to complete the 3-way handshake. At this point, a two-way communication application may be started at 490.

Once a PATH message is sent out from the sender 310, the message travels through 15 the network, from edge policy enforcement device to edge policy enforcement device, before it reaches the receiver 320. FIG. 5 shows the flowchart of another embodiment of the invention for an ingress policy enforcement device. In FIG. 5, upon intercepting a PATH message at an ingress policy enforcement device at 510, the received PATH message is processed at 520. Using the information carried in the received PATH message, the ingress 20 policy enforcement device reserves the network resource.

The reservation may be made by communicating with the PDP in the same domain as a policy enforcement device using COPS-RSVP. By examining the reservation request from the policy enforcement device against the available resources and the network policies, the

PDP decides whether the request will be granted or not. The decision is then communicated back to the policy enforcement device. A different embodiment for reserving network resource is directly through the policy enforcement device without consulting with the PDP. In this case, the admission required domain wide information is available to the policy enforcement device and the policy enforcement device is entrusted by the network administrators to make resource reservation decisions based on local knowledge.

In Fig. 5, whether the reservation is made through the PDP is determined at act 530. If the reservation needs to be made through the PDP, the policy enforcement device communicates with the PDP at act 535 using COPS-RSVP. If the policy enforcement device can reserve resource directly, the resource is reserved at act 537 directly by the policy enforcement device. The reservation may succeed or fail, depending on, for example, the availability of the network resources, the admission policies, as well as the resources that are needed.

If the reservation is not successful, determined at act 540, the ingress policy enforcement device may construct a PATH_ERR message at 550 and sends it back at act 560 to the sender 310 to inform an unsuccessful reservation for the forward direction. If the reservation is successful, the ingress policy enforcement device forwards the PATH message to an egress policy enforcement device of the same domain at 570.

As indicated in FIG. 3, an egress policy enforcement device receives and processes messages in both first and second passes. A PATH message is intercepted by an egress policy enforcement device in the first pass and an RESV message is passed to an egress policy enforcement device in the second pass. FIG. 6 is a flowchart for an egress policy enforcement device. When a message is received at an egress policy enforcement device at 605, it is

- examined to see whether it is a PATH message or an RESV message. If the received message is a PATH message, the egress policy enforcement device processes the PATH message at 615 and adds its own address to the NEXT_HOP object of the PATH message at 620, making sure that the RESV message in the second pass will be sent to this policy enforcement device.
- 5 The egress policy enforcement device determines an ingress policy enforcement device of the next domain at 635 and forwards the PATH message to the ingress policy enforcement device at 630. The egress policy enforcement device then returns to 605 to wait for the arrival of next message.

When the message received by an egress policy enforcement device is an RESV

- 10 message (decided at 610), it indicates that the reservation for the forward direction has been successful. This RESV message, carrying the reservation information for the reverse direction, initiates the resource reservation for the reverse traffic. The egress policy enforcement device processes the RESV message at 640.

Based on the reservation information carried in the RESV message, the egress policy

- 15 enforcement device determines, at act 643, whether the needed network resource needs to be reserved through the PDP. If the reservation is to be made through the PDP, the policy enforcement device consults with its PDP at 645 and receives a decision from the PDP. If the policy enforcement device can make reservation directly, the resource is reserved at act 647.

If the reservation is successful, determined at act 650, the egress policy enforcement

- 20 device forwards the RESV message to the next egress policy enforcement device at 670 using the addresses in the NEXT_HOP object of the RESV message. If the reservation is not successful (the required resources are not granted), the egress policy enforcement device constructs error messages and sends to both the sender 310 and the receiver 320. At act 660,

an RESV_ERR message is constructed and sent, at 665 to the receiver 320, signaling that the reservation request initiated by the receiver 320 has failed. The egress policy enforcement device may also construct a PATH_ERR message, at act 670, and send it, at act 675, to the sender 310 to indicate a failure in reserving the network resource in the reverse direction.

- 5 Resource reservation in either direction is performed via the communication between an edge policy enforcement device (ingress policy enforcement device in the forward direction and egress policy enforcement device in the reverse direction) and the PDP of the same domain. FIG. 7 illustrates a flowchart for a PDP. Upon receiving a reservation request at 710, the PDP processes the request at 720 and checks with the network policies as well as the available resources of its corresponding domain at 730. If the policies allow and the requested resources are available, the PDP may decide to admit the request by installing necessary per-flow filters as well as traffic profiles at 760 via COPS-PR. The PDP then issues an admission at 770 to the requesting policy enforcement device. If the request is not granted, the PDP informs the requesting policy enforcement device its decision at 750. A successful consultation between a requesting policy enforcement device and a PDP results in required network resource being reserved at the corresponding network domain. Similar to the reservation for forward direction traffic in the first pass, when an RESV message travels hop-to-hop, network resource required for the traffic in the reverse direction are reserved at each stop. When the RESV message reaches the sender 310, the resources required for the reverse direction have been reserved along the path.
- 10 15 20

FIG. 8 and FIG. 9 show the flowchart for the receiver 320. Once the sender 310 initiates a 3-way handshake, if the receiver 320 receives a PATH message at 810, it indicates that the resource reservation for the forward direction is successful. The receiver 320 responds

to the PATH message and, at the same time, initiates the reservation for the reverse direction by constructing an RESV message. To do so, the PATH message is processed at 820. The PATH message has an NEXT_HOP object that contains the route information from the sender 310 to the receiver 320. Such information may be extracted at 830 and used to construct an 5 RESV message at 840. The RESV message generated by the receiver 320 carries both the reservation information for the reverse direction as well as the path information instructing how the RESV message should travel from the receiver 320 to the sender 310 (NEXT_HOP). The RESV message is sent from the receiver 320 to the egress policy enforcement device of the last domain between the sender 310 and the receiver 320. The receiver 320 marks the time 10 853 to establish the time reference to be used in a time-out mechanism and then waits for return messages.

If a return message is not received before a time-out at 855, the handshake is aborted. If a return message is received within the time-out at 855 in FIG. 9, it is first examined at 860 to see whether it is an RESV_ERR message. If the received message is an RESV_ERR, it 15 means that the resource reservation for the reverse direction has failed. In this case, the receiver 320 aborts the 3-way handshake. If the received message is an RESV_Confirm message, it signals a successful 3-way handshake, meaning that the reservation for both directions has succeeded. In this case, the receiver 320 enters a two-way communication session at 870. If the received message is neither an RESV_ERR nor an RESV_Confirm, 20 processing returns to 855 to await the next message.

The embodiment of the invention illustrated in FIG. 3 is a sender-driven 3-way handshake RSVP scheme that reserves network resources for both forward and reverse directions. A different embodiment of the present invention is a receiver-driven 4-way

handshake RSVP scheme that reserves needed network resources in both forward and reverse directions and that supports multicast applications. FIG. 10 illustrates this embodiment of the invention, in which the two parties are 1010 (the sender) and 1020 (the receiver). Similar to the exemplary illustration in FIG. 3, there are four domains in FIG. 10. E₁, E₃, E₅, and E₇ (1030a, 1030b, 1030c, 1030d) are the ingress policy enforcement devices and E₂, E₄, E₆, and E₈ (1040a, 1040b, 1040c, 1040d) are the egress policy enforcement devices of the four illustrated network domains, looking in the direction from the sender 1010 to the receiver 1020. The PDPs for the four domains are 1050a, 1050b, 1050c, and 1050d. Both ingress and egress policy enforcement devices may communicate with their domain PDPs via COPS-
5 10 RSVP to perform resource reservation.

In FIG. 10, the sender 1010 initiates a 4-way handshake. A first PATH message travels through the network, in the first pass, to probe a path between the sender 1010 and the receiver 1020. The resources needed in the forward direction are reserved in the second pass (initiated or driven by the receiver 1020). The resources needed in the reverse direction are
10 15 reserved in the third pass and the reservations in the forward direction are confirmed. The last pass finishes the 4-way handshake by confirming the reservation in the reverse direction.

To start a 4-way handshake, the sender 1010 generates a first PATH message, PATH₁. Message PATH₁ does not contain reservation information. It is for probing a path from the sender 1010 to the receiver 1020. When PATH₁ travels from policy enforcement device to
20 policy enforcement device, the path is recorded by adding the address of each visited edge policy enforcement device to the NEXT_HOP object of PATH₁. When PATH₁ arrives at the receiver 1020, a first RESV message, RESV₁ responding to PATH₁, is generated at the receiver 1020. RESV₁ carries the reservation requests for the forward traffic. At the same

time, RESV₁ carries a PATH object, PATH₂, that serves as a second PATH message for the reverse direction.

The coupled RESV₁+PATH₂ message travels hop-by-hop to the edge policy enforcement devices of each domain. Each ingress policy enforcement device (1030a, 1030b,

- 5 1030c, 1030d) intercepts the RESV₁ message and performs policy and bandwidth admission control by communicating with the PDP of the same domain (1050a, 1050b, 1050c, or 1050d) using COPS-RSVP. If the request is granted, the RESV₁+PATH₂ message is forwarded to the egress policy enforcement device (1040a, 1040b, 1040c) of the next domain to continue the reservation request along the path. When the RESV₁+PATH₂ message reaches the egress 10
10 policy enforcement device, it ensures that its address remains as part of the NEXT_HOP object. This ensures that the RESV₂ message will eventually be sent to it.

When message RESV₁/PATH₂ reaches the sender 1010, a receiver-driven reservation for the forward direction is completed. In this case, the sender 1010 generates a message including an RESV₁_Confirm portion, as a response to RESV₁, and an RESV₂ portion, as a 15 response to PATH₂. The former is to acknowledge the successful reservation in the forward direction while the latter is to initiate the resource reservation for the reverse direction. The sender 1010 sends the coupled RESV₂+RESV₁_Confirm message to the receiver 1020 via all the egress policy enforcement devices along the path. At each egress policy enforcement device, based on the reservation request carried by RESV₂, policy and bandwidth admission

- 20 control is performed by consulting with the PDP of the corresponding domain using COPS-RSVP. If the reservation is admitted, the PDP installs necessary per-flow filters and traffic profiles via COPS-PR. If message RESV₂+ RESV₁_Confirm successfully reaches the receiver 1020, the reservation in both directions is completed. The receiver 1020 then generates a

RESV₂_Confirm message and sends it directly to the sender 1010 to complete the 4-way handshake.

In the embodiment illustrated in FIG. 10, a two-way reservation is considered successful only when the admission control decisions are granted in both forward and reverse directions. A failure in reserving the network resources needed in the forward direction at any ingress policy enforcement device may be signaled by sending an RESV_ERR message from that policy enforcement device to the receiver 1020. A failure in reserving the network resource needed in the reverse direction at any egress policy enforcement device may be signaled by sending an RESV_ERR message back to the sender 1010.

FIG. 11-13 show a flowchart for the sender 1010. To initiate a 4-way handshake, the sender 1010 constructs the first PATH message, PATH₁, at 1105. Message PATH₁ is sent at 1110 and the time may be marked at 1115 so that a first reference time for a time-out mechanism may be established. The sender 1010 then waits for a return message.

If a message is received before a time-out at 1120, the message type is determined at act 1130. If it is not an RESV₁+PATH₂ message, the sender 1010 goes back to 1120 to wait. If the time-out condition is satisfied at 1120, the sender 1010 aborts the 4-way handshake. A time-out tells the sender 1010 that the forward direction reservation failed and an RESV₁_ERR message will be sent, in this case, to the receiver 1020.

If the received message is an RESV₁+PATH₂, it indicates that the reservation in the forward direction has been successful. In this case, the sender 1010 constructs a third message, a coupled RESV₁_Confirm+RESV₂ message. The former is to acknowledge received RESV₁ and the latter is to initiate the reservation for the reverse direction. An RESV₁_Confirm message is generated at 1145 of FIG. 12. At the same time, the sender 1010

constructs a RESV₂ message. To do so, the PATH₂ message is processed at 1150. The NEXT_HOP object of PATH₂ is constructed in the second pass, during which egress policy enforcement devices add their addresses to the NEXT_HOP object of PATH₂ so that the path represented by the NEXT_HOP object contains only the egress policy enforcement devices.

- 5 The NEXT_HOP from PATH₂ is used to construct the RESV₂ message at 1160. The coupled RESV₁_Confirm+RESV₂ message is sent to the receiver 1020 at 1170, traveling through only egress policy enforcement devices in the forward direction. The sender 1010 then waits for either an error message RESV₂_ERR, informing the sender 1010 that the reservation for the reverse direction fails, or a confirmation message RESV₂_Confirm,
10 indicating that the reservation for the reverse direction succeeds.

- The sender 1010 intercepts a message at act 1173. If the message is an RESV₂_ERR, determined at 1175, the sender 1010 aborts the 4-way handshake at 1195. If the message is an RESV₂_Confirm, determined at act 1180, it means that the reservation for both directions (forward and reverse) has been successful and the four-way handshake is complete. The
15 corresponding two-way communication, in this case, may be started at 1185.

- In the RSVP reservation scheme shown in FIG. 10, an ingress policy enforcement device performs different functions, depending on the type of message it receives. In the first pass, an ingress policy enforcement device receives a PATH₁ message. In the second pass, an ingress policy enforcement device receives a coupled RESV₁+PATH₂ message. FIG. 14
20 shows the flowchart for an ingress policy enforcement device in a four-way handshake scheme.

Upon intercepting a message at an ingress policy enforcement device at 1410, the message type is examined at 1415. If the received message is a PATH₁ message, it is in the

first pass of the 4-way handshake. In this case, the ingress policy enforcement device processes PATH₁ and add its own address to the NEXT_HOP object of PATH₁ at 1425 that ensures that the RESV₁+PATH₂ message will be sent to this ingress policy enforcement device. The revised PATH₁ message is then forwarded toward the egress policy enforcement device of the same network domain at 1430. The ingress policy enforcement device then returns to a receiving mode at 1410.

If the received message is a RESV₁+PATH₂ message, it is in the second pass of the 4-way handshake. In this pass, an ingress policy enforcement device performs both the function of reserving resources needed for the forward direction (based on the RESV₁ message) and

10 the function of processing the second PATH message, PATH₂, for the reverse direction.

To reserve requested resources for the forward direction, the ingress policy enforcement device processes RESV₁ message at 1440. Whether the resources are to be reserved through the PDP is determined at act 1443. If the reservation is to be made through the PDP, the policy enforcement device communicates with the PDP of the same domain at

15 1445. The communication may be performed using protocol COPS-RSVP. Based on available resources and the network policies, the PDP decides whether the resource request for the forward direction will be granted or not. Such a decision is communicated back to the ingress policy enforcement device.

The policy enforcement device may also reserve the resource directly at act 1447. If 20 the reservation is successful, determined at act 1450, the ingress policy enforcement device forwards the received RESV₁+PATH₂ message at 1470 to the egress policy enforcement device of the next domain in the reverse direction. If the request for the resources required in forward direction is not granted (at 1450), the reservation fails. In this case, the ingress policy

enforcement device constructs an error message RESV₁_ERR at 1455 and sends it at 1460 to the receiver 1020, signaling a failure in reserving required resources in the forward direction.

As described in FIG. 10, an egress policy enforcement device in a 4-way handshake may receive three types of messages. A PATH₁ message passes through egress policy enforcement devices in the first pass, an RESV₁+PATH₂ message passes through egress policy enforcement devices in the second pass, and an RESV₂+RESV₁_Confirm message passes through egress policy enforcement devices in the third pass. Depending on the type of the received message, an egress policy enforcement device performs different functions. FIG. 15 presents a flowchart for an egress policy enforcement device in a four-way handshake scheme.

When a message is received at an egress policy enforcement device at 1510, its type is examined at 1515. As depicted in FIG. 15, if the received message is a PATH₁ message, the egress policy enforcement device processes the received PATH₁ message at 1520 and adds its own address to the NEXT_HOP object of PATH₁ at 1525. The revised PATH₁ is sent to the ingress policy enforcement device of the next domain in forward direction at 1535. The egress policy enforcement device then waits to receive the next message at 1510.

If the message received at 1510 is an RESV₁+PATH₂ message, the egress policy enforcement device adds its own address to the NEXT_HOP object of PATH₂ at 1536 and forwards the RESV₁+PATH₂ message at 1537 to the ingress policy enforcement device of the same domain (in the reverse direction). The egress policy enforcement device then goes back to a waiting mode at 1510 to intercept the next message.

If the message received at 1510 is a coupled RESV₂+RESV₁_Confirm message, it indicates that the resources needed for the forward direction have been successfully reserved

and the reservation for the reverse direction needs to be made. The received RESV₂ message is originated at the sender 1010, carrying the reservation request for the reverse direction. RESV₂ message is processed at 1540.

Based on the reservation request in RESV₂, the egress policy enforcement device
5 reserves needed resource either through the PDP at act 1545 or directly at act 1547, depending
on the decision made at act 1543 in terms of how the resource is to be reserved. If the
reservation is successful, determined at act 1550, the egress policy enforcement device
forwards the RESV₂+RESV₁_Confirm message at 1570 to the next egress policy enforcement
device, using the addresses defined in the NEXT_HOP object of RESV₂ message. If the
10 reservation request is not granted, the egress policy enforcement device constructs a
RESV₂_ERR message at 1555 and sends it at 1560 back to the sender 1010, informing the
receiver 1020 that the reservation request initiated by the receiver 1020 in the reverse
direction has failed.

FIG. 16 and FIG. 17 show the flowchart for the receiver 1020 in a 4-way handshake
15 scheme. Once a 4-way handshake is initiated by the sender 1010, when the receiver 1020
receives a PATH₁ message at 1610, the NEXT_HOP object of PATH₁ message contains the
addresses of all the edge policy enforcement devices that PATH₁ travels through. The
NEXT_HOP object defines a path between the sender 1010 and the receiver 1020 and this
path is used to construct a RESV₁ message at 1620. The RESV₁ message carries also the
20 reservation request for the forward direction and travels along the path defined by the
NEXT_HOP from PATH₁.

The receiver 1020, at the same time, also constructs a PATH₂ message that serves as
the PATH message of the resource reservation for the reverse direction. The receiver 1020

then sends the coupled RESV₁+PATH₂ message at 1630 to the first egress policy enforcement device in the reverse direction. This starts the second pass of the 4-way handshake. Message RESV₁+PATH₂ travels hop-by-hop to all the edge policy enforcement devices in the reverse direction. Along this reverse path, the reservation for the forward direction may be made,

5 based on RESV₁, at each of the ingress policy enforcement devices. Also in this second pass, the reverse path defined by egress policy enforcement devices only, is constructed, along the way, and the addresses of the egress policy enforcement devices in the reverse path are recorded in PATH₂ message as the NEXT_HOP object of PATH₂ and it provides the path for the third pass of the 4-way handshake, in which the resource reservation for the reverse

10 direction will be made.

Once the RESV₁+PATH₂ message is passed on, the receiver 1020 enters a waiting mode for return messages. The receiver 1020 may mark the time (at 1633) so that a reference time for a time-out mechanism may be established.

A return message received at the receiver 1020 may be an RESV₁_ERR message,

15 which indicates that the reservation for the reverse direction has failed, or a coupled RESV₁_Confirm/RESV₂ message, which indicates that the reservation in both directions has succeeded. Depending on the type of message received the receiver 1020 functions differently, as illustrated in FIG. 17. If a message is not received before time-out at 1635, the 4-way handshake is aborted at 1670. If a message is received and the received message is an

20 RESV₂_ERR message at 1640, the receiver 1020 also aborts the 4-way handshake at 1670. If the timely received message is an RESV₂/RESV₁_Confirm message at 1645, the receiver 1020 constructs an acknowledgement message RESV₂_Confirm at 1690 and sends it directly

back to the sender 1010 at 1655 to complete the 4-way handshake. A 2-way communication may then be started at 1660.

The processing described above may be performed by a general-purpose computer alone or in connection with a special purpose computer. Such processing may be performed

5 by a single platform or by a distributed processing platform. In addition, such processing and functionality can be implemented in the form of special purpose hardware or in the form of software being run by a general-purpose computer. Any data handled in such processing or created as a result of such processing can be stored in any memory as is conventional in the art. By way of example, such data may be stored in a temporary memory, such as in the

10 RAM of a given computer system or subsystem. In addition, or in the alternative, such data may be stored in longer-term storage devices, for example, magnetic disks, rewritable optical disks, and so on. For purposes of the disclosure herein, a computer-readable media may comprise any form of data storage mechanism, including such existing memory technologies as well as hardware or circuit representations of such structures and of such data.

15 While the invention has been described with reference to the certain illustrated embodiments, the words that have been used herein are words of description, rather than words of limitation. Changes may be made, within the purview of the appended claims, without departing from the scope and spirit of the invention in its aspects. Although the invention has been described herein with reference to particular structures, acts, and materials,

20 the invention is not to be limited to the particulars disclosed, but rather extends to all equivalent structures, acts, and, materials, such as are within the scope of the appended claims.